txture

Cloud Deep Dive Series - Part 1

SURVIVING IN THE MESSY WORLD OF MULTI-CLOUD PRICING

txture

Surviving in the Messy World of Multi-Cloud Pricing

CONTENT

EXECUTIVE SUMMARY 03 — 1 min

INTRODUCTION 04 — 2 min

MAJOR CHALLENGES 05-09 — 8 min

TXTURE SOLUTION APPROACH 10-11 — 3 min

CONCLUSION 12 — 1 min

TXTURE PLATFOM 13 — 2 min

EXECUTIVE SUMMARY

SURVIVING IN THE MESSY WORLD OF MULTI-CLOUD PRICING

1 Minute

The comparability between cloud providers and the predictability of cloud costs are major concerns for organizations considering the move to the cloud. Those uncertainties about the costs accumulating for running a large portfolio of applications in the cloud can be a major hurdle when migrating on-premises solutions.

In this first installment of our series “Cloud Deep Dives”, we want to take a closer look at the issue based on some examples. Here we aim to bring clarity to the challenges of the cloud pricing matter. Furthermore, we share some of our practical experience and approaches to make Multi-Cloud pricing more transparent - specifically for Cloud Center of Excellence (CCoE) teams.

In this article you are going to read about the following aspects of Multi-Cloud pricing:

- What causes the lack of comparability of cloud products and complexity in Multi-Cloud pricing?
- What are the five major challenges in Multi-Cloud pricing organizations are facing during their cloud transformation?
- How can you master these challenges and identify a cost-optimized cloud product portfolio?

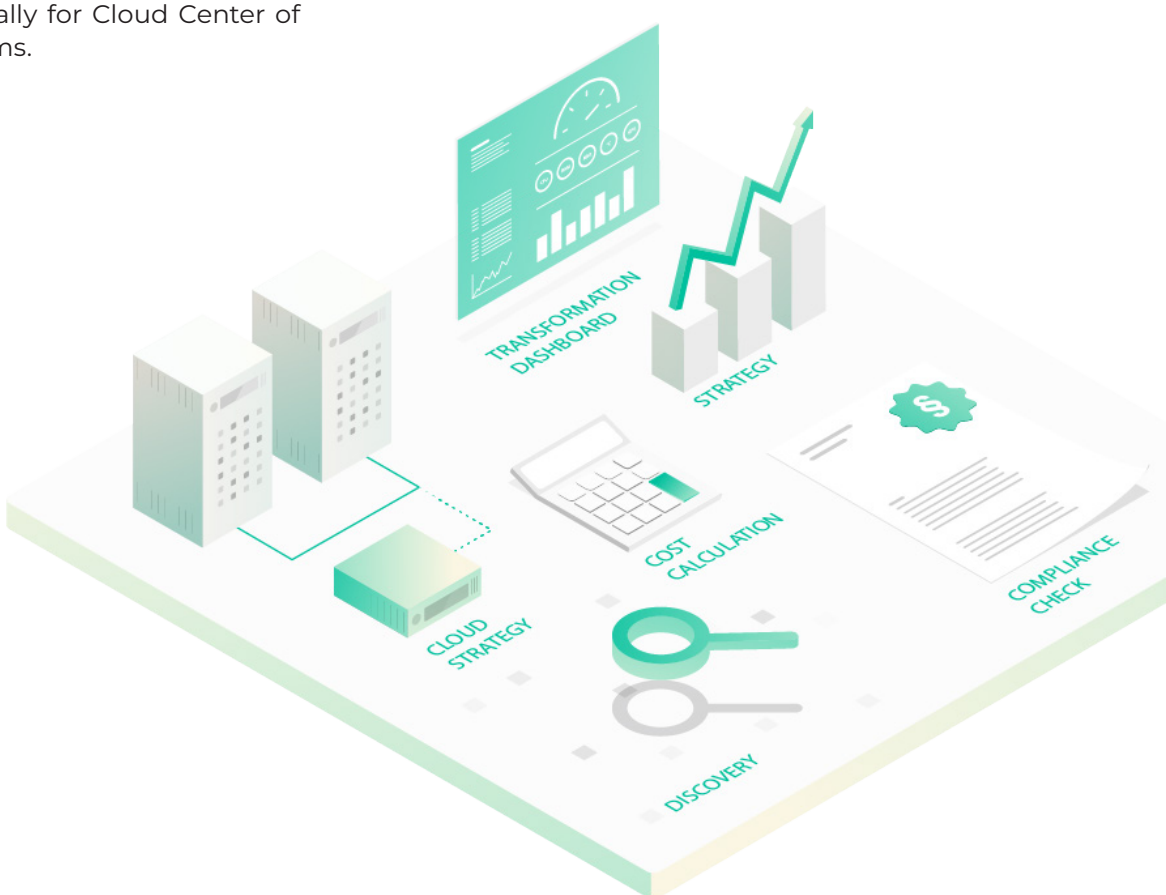

INTRODUCTION

COMPARABILITY AND PREDICTABILITY OF CLOUD COSTS

2 Minutes

The current situation on the cloud market regarding pricing politics can be described as confusing or even obscure. Pricing models are complex and especially if cloud services from several vendors are to be compared or even combined (like it is the case for Multi-Cloud), cost transparency becomes very difficult to achieve.

Especially for globally operating organizations with hundreds of applications, the pricing issue has several mutually reinforcing characteristics. These includes the wide diversity of cloud services to choose from, pricing fluctuations, difficulties in cost prediction or a lack of comparability of services from different vendors. Furthermore, complexity is also reinforced by the size of business applications and the number of related technical components.

To handle all this, leading companies have started to establish workforces dedicated exclusively to drive their cloud transformation. This article is particularly relevant for this new type of organizational institution often known as Cloud Center of Excellence (CCOE) or Cloud Competence Center (CCC). Since one of their key roles is to give advice and develop best practices for cloud cost forecasting, their daily business is to survive in the messy world of Multi-Cloud pricing.

PROBLEM DEFINITION

Before we can start, we need to define the term “Multi-Cloud” to ensure a consistent understanding. Gartner defines Multi-Cloud as the use of multiple cloud computing vendors at once[1]. The focus of this term is that products from multiple providers are used for one heterogeneous IT architecture[2].

The overall pain of cloud product pricing complexity in a Multi-Cloud context can be derived from the following two questions:

1. How can we achieve price comparability across similar products from different vendors to achieve the best value for our Multi-Cloud product portfolio?

2. How can we achieve cost transparency and predictability for our Multi-Cloud product portfolio?

These two questions are hard to answer due to the following five major challenges in Multi-Cloud pricing explained in the following chapter.

1 https://www.gartner.com/smarterwithgartner/modernize-it-infrastructure-in-a-hybrid-world/
2 https://www.gartner.com/imagesrv/media-products/pdf/telefonica/Telefonica-1-4C6KXTB.pdf

MAJOR CHALLENGES

CHALLENGE 1 - CLOUD SERVICE DIVERSITY

8 Minutes

The first and probably most critical challenge arises from the tremendous diversity of cloud products, product variants, and pricing options. While the number of cloud providers seems to be still traceable (at least if we just focus on globally operating hyperscalers), it is easy to lose track of a provider's product portfolio. The comparison between providers considering their product configurations, service levels, and pricing options is even more challenging.

In a nutshell, the complexity of calculating cloud product pricing can be derived from five factors:

1. The variety of cloud vendors.
2. The variety of products provided by a single or various cloud vendors.
3. The variety of product variants, for example, the location, instance size or product generation. Some providers even allow to adjust their products variants' scope of service based on licenses. This makes it possible to buy and add single features to a product variant.
4. The variety of parameters used for pricing calculations, for example, runtime per month, required storage, amount of transactions or support included.
5. The variety of billing options, such as on-demand prices, discounts for reserved instances, volume discounts or upfront costs.

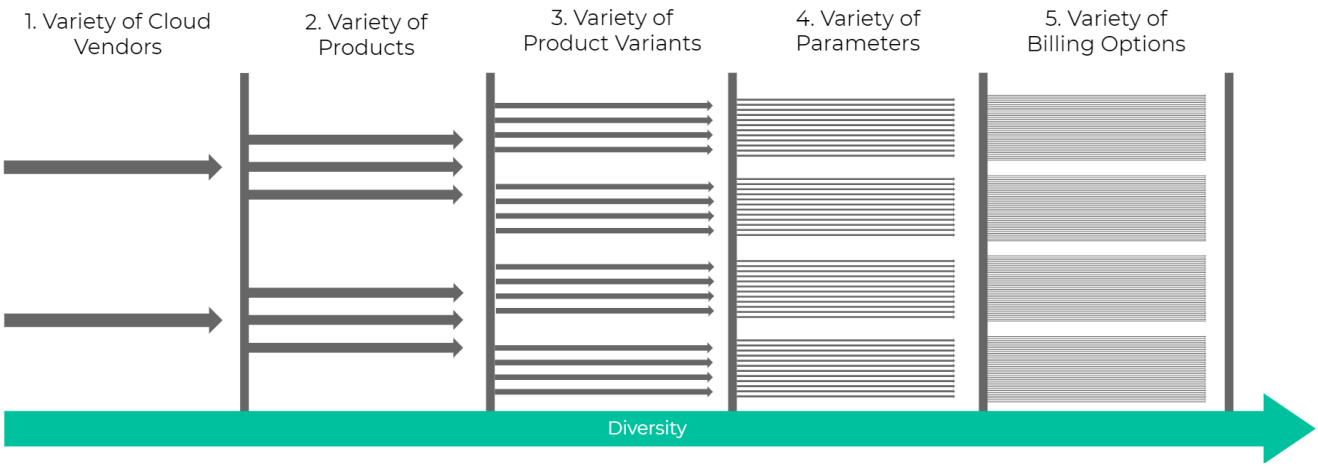

Figure 1: Multipliers of pricing diversity

EXAMPLES

We would like to illustrate the complexity of comparing products and prices with two practical examples. To do so, we browsed through the product catalog of two well-known cloud providers, Microsoft Azure and Amazon Web-services (AWS), and looked for equivalent products with similar configurations.

Both AWS and Azure offer pricing calculators, where a product has to be selected first, followed by a variety of product and billing configurations that strongly impact the final pricing. To keep it as simple as possible, we chose two basic products - the "Azure Virtual Machine" versus "Amazon EC2".

We selected instances that are as similar as possible. However, for both providers there were a few things that made this more complex and had to be considered:

In the case of Azure, we need to think about which Pricing Tier and what Series we want to choose a VM from. Currently, there are more than ten different VM Series available. Each Series has its strengths and weaknesses and is suitable for different use cases. The A-Series, for example, is a low-priced option that offers instances that are suitable for development or test servers. In contrast, the M-Series has completely different capabilities. It offers the largest memory-optimized virtual machines and can be up to a hundred times more expensive than A-Series instances.

Moreover, Azure offers a variety of Pricing Tiers for every product where each is suitable for another use case. There are the Tiers Free, Basic, Standard, and Storage Optimized. The Free Tier, for example, is typically used to evaluate a service. Basic can be used for a development and production deployment and the "Standard" Tier for production[3]. The Tier has an impact on the service limit of a product like limits on storage, workloads, and quantities of indexes, documents, and other objects[4].

In the AWS world, VM instances (EC2) are categorized in use cases and Instance Types instead of Azure's Series and Tiers. Instance Types are assigned to use cases, such as "Compute Optimized", "Memory Optimized" or "Storage Optimized". They comprise varying combinations of CPU, memory, storage, and networking capacity. For each Instance Type, a variety of sizes are offered. Of course, AWS provides a wide variety of different instances with varying capabilities too.

Azure	AWS

Azure Virtual Machines - D-Series
· Region: Central Germany
· Tier: Standard
· Series, Size & Version: D4 v3
  · Processors: 4 vCPUs
  · RAM: 16 GB
· OS: Windows (Type: OS only)
· Runtime: 730h/Month
· Billing Option: Pay as you go
	Amazon EC2 - M5
	· Region: EU (Frankfurt)
	· Category: General Purpose
	· Instance Type & Size: m5.xlarge
	  · Processors: 4 vCPUs
	  · RAM: 16 GB (approx 17 GB)
	· OS Windows Server
	· Runtime: 730h/Month
	· Billing Option: On-demand

Azure Virtual Machines B-Series
· Region: Central Gemany
· Tier: Standard
· Series, Sive & Version: B4MS
  · Processors: 4 vCPUs
  · RAM: 16GB
· OS: Windows (Type: OS only)
· Runtime: 730h/Month
· Billing Option: Pay as you go
	Amazon EC2 - T3a
	· Region: EU (Frankfurt)
	· Category: General Purpose
	· Instance Type & Size: t3a.xlarge
	  · Processors: 4 vCPUs
	  · RAM: 16 GB (approx 17 GB)
	· OS: Windows Server
	· Runtime: 730h/Month
	· Billing Option: On-demand

Table 1: Comparison of Azure and AWS Cloud Products

3 https://docs.microsoft.com/en-us/azure/search/search-sku-tier
4 https://docs.microsoft.com/bs-latn-ba/azure/search/search-limits-quotas-capacity

For the first comparison, we chose a general-purpose VM with 4 vCPUs, approx. 16GB of RAM and Windows OS hosted in central Germany.

Example 1: D-Series vs. M5 (prices from Nov 7, 2019) At Azure, an appropriate D-Series VM instance cost $334.34 for one month at that time. A similar configuration for an M5 AWS instance costs only $302.22.

For both products, we chose a billing option which gave us full flexibility without the need to commit for a period or any upfront payments. Azure calls this "Pay as you go" whereas for AWS it is the "On-Demand" pricing model. If we had accounted for different billing options or other parameters such as transaction costs, the price could have changed radically.

Since matching AWS and Azure products are not always clear, let us compare one more VM pair.

Example 2: B-Series vs. T3a (prices from Nov 7, 2019)The VMs of the Azure B-Series or the AWS T3a Instance Type are cheap instances that can be used if a consistently high CPU load is not needed but the coverage of some peaks per hour is sufficient. These instances allow only a few peaks per hour, however, they are much cheaper. With the same configurations as before, such an instance at Azure costs $104.94. An alternative AWS instance costs $179.87.

Azure significantly wins the second comparison at first glance. However, we did not consider an important factor. We do not know how long each instance can burst to maximum power per hour. That makes comparison quite difficult. Although both providers, AWS and Azure, greatly improved the usability of their pricing calculators, in such complex situations difficulties in comparing products cannot be resolved by a calculator that easily.

MAJOR CHALLENGES

CHALLENGE 2 - PRICE AND SERVICE FLUCTUATION

The second challenge is directly related to the first one. Prices change rapidly over time and new products are released while older ones disappear from the vendors' product portfolio. Fluctuations of most providers' product portfolios are therefore the next issue regarding the comparability of cloud product prices. AWS, for example, currently (November 2019) offers 169 products (counted based on namespaces).

In 2018 there have been only 95 products. For other providers, the situation is similar. For each product there exist various product variants. For each variant, you can choose from a variety of billing options. While the range of IaaS products and related prices are still manageable, the diversity for PaaS is already much higher. In the SaaS area, diversity gets nearly infinite.

Thus, changing prices and product portfolios create another dimension of complexity.

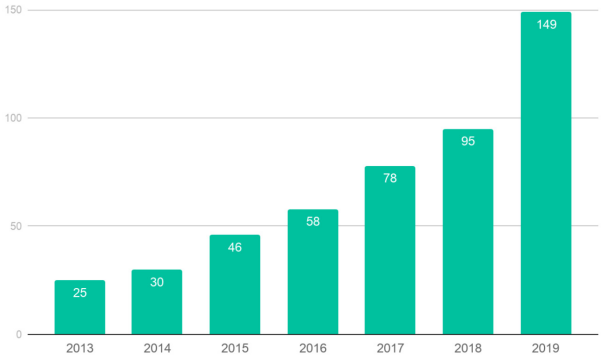

Figure 2: Number of Services (Y-axis) in the AWS Service Cloud over time (based on namespaces)[5]

5 https://medium.com/cloudpegboard/how-many-aws-services-are-there-51dda44fa946

MAJOR CHALLENGES

CHALLENGE 3 - COST PREDICTABILITY

The next challenge in comparing and estimating prices for a potential cloud solution arises from a key strength of cloud computing - Scalability. Scalability allows one to quickly increase the size or computing power of a cloud service and decrease it if not needed anymore. It allows us to easily cover peak loads and save resources. Consequently, the utilization of a system and associated costs can be optimized.

The problem resulting from cloud pricing complexity issues is the loss of predictability. In on-premises situations, it has been much easier to predict costs for acquiring and operating your own IT Infrastructure than predicting required loads for the next months or even years.

While CAPEX was the predominant cost factor for running an on-premises IT infrastructure, cloud service costs are driven solely by OPEX. Also, there are hardly any upfront costs for operating a cloud portfolio. This paradigm shift fundamentally changes an organization's IT budget plan[6]. If no historical data or experiences exist for estimating future loads, challenge 3 hardly allows the plannability of cloud costs.

MAJOR CHALLENGES

CHALLENGE 4 - SIZE AND SERVICE MODEL OF BUSINESS APPLICATIONS

When having a detailed look at the infrastruc-ture level of large business applications, the interweaving of a high number of individual and interdependent components becomes visible. In extreme situations, such applications consist of several hundred components, such as interfaces, databases, load balancers, middleware or different types of databases and VMs. Migrating such an application requires consideration of these components on both a technical and an economic level.

The analysis of the IT infrastructure is also an important factor when deciding on the cloud service model. Costs strongly differ by whether the application is migrated to IaaS, PaaS or CaaS. Depending on the service model, suitable cloud product replacements need to be found for a high number of on-premises components. Thus, already estimating the cloud cost of a single application can be challenging.

6 James Bond, The Enterprise Cloud: Best Practices for Transforming Legacy IT (Sebastopol, CA: O'Reilly, 2015), pp. 200-201.

MAJOR CHALLENGES

CHALLENGE 5 - AUTOMATION

The perception of the first four challenges seems to be that to identify a cost-optimized composition of cloud products manually is hardly possible without great effort, especially if we want to go towards multiple providers. The logical consequence is an automated approach for cost estimation.

Cloud vendors offer calculators that allow to semi-automatically estimate the prices for a specific set of cloud products. However, they only take into account their own products. Comparability between different vendors is not supported. Tool vendors who offer automated solutions for cloud cost optimization are facing difficulties in programmatically accessing live pricing data from cloud vendors. Some of the major vendors provide APIs, some do not. Even though APIs make it easier to access live data, in many cases data is inconsistent and first has to be prepared before it can be used for further calculations. In the following example, multiple prices need to be combined before we get a result for the corresponding cloud product.

EXAMPLE

Let us calculate the price for a simple example, Google's Compute Engine. Via the API we can access live-prices for a single vCPU of a specific type in a specific location. Moreover, we can access prices per GB of memory.
We would like to have a N1 standard machine. As region, we choose Zürich (europe-west6). From the machine-type API we get the configuration of N1 standard machine instance: It consists of 4 vCPUS (number of CPUs = 4) and 15 GB of memory (RAM in GB = 15). The price per vCPU for this type in that region is currently at $0.044231 / vCPU hour (November 7, 2019). The memory price is $0.005928 / GB hour. Both prices we got from a pricing API.

Since we know what components our Google Compute Engine instance consists of and we know the prices for all single components, we can calculate the price for this product. However, we have to do this manually. Later we can automatically keep these prices up-to-date via APIs. The price per hour for our N1 standard machine is $0,2658 per hour.

First calculate prices for CPU and RAM separately:

```
Number of CPUs * price per CPU =
4 * 0.044231 = 0.176924
RAM in GB * price per GB =
15 * 0.005928 = 0.08892
```

Now we can combine prices to get the instance price per hour:

```
0.08892 + 0.176924 = 0.265844
(~0.2658)
```

The quality of data that can be accessed via API differs from vendor to vendor. Often, however, the data must first be consolidated as shown in the example. That can be very time consuming and cumbersome.

TXTURE SOLUTION APPROACH

BEST PRACTICE FOR COMPARING MULTI-CLOUD PRICES

3 Minutes

Now that we have defined the five major challenges in estimating and comparing Multi-Cloud prices, we want to share our approach to handle the problem of cloud pricing calculation. Basically, our approach consists of three steps.

1. AUTOMATIC DATA COLLECTION VIA APIS AND QUALITY ASSURANCE

First, we implemented interfaces to all major cloud providers. We access most data from the providers' product catalogs via REST APIs. For some products, it is indispensable to retrieve the data manually from the vendor. For quality assurance we have a dedicated team that constantly validates the pricing data.

As seen in Example 3, in many cases, product prices must be normalized first. The result is an up-to-date and quality-assured dataset of cloud products and their variants including normalized prices. In addition to prices, this dataset also includes important information about each product variant's properties like the size, the number of CPUs or RAM-size. Furthermore, we can use pricing information to identify which of a provider's product variants are available in what region. If we get a price for a product variant for a specific region, we can assume that this product is available in that region. In addition to prices, we can now automatically track the availability of product variants for each cloud provider's region.

At this point, however, prices and products cannot be compared yet. That is why we need to bring all prices to a common denominator which is done in the second step.

2. MAPPING DATA ONTO AN OVERARCHING DATA MODEL

Different cloud providers have entirely different pricing schemes for similar services. To compare the price, we need one model which can deal with both pricing schemes. That is why, we established an overarching pricing data model. Prices for all products and product variants are matched onto this data model. We use a unified concepts to model runtime costs, network costs or costs for storage. Due to the very large amount of different product variants, this step is unavoidable. Although some details are thus no longer taken into account, for large applications a cost estimation is still very helpful.

3. ESTIMATION OF PRICES FOR PRODUCT SUGGESTIONS

Now that we have our pricing dataset which is continuously updated automatically as explained in the previous steps, we can start to estimate costs for our desired cloud product portfolio.

All further decisions about the composition of a cloud portfolio are derived from the existing on-premises IT infrastructure. Since it is not the focus of this article, I do not go into detail about IT infrastructure assessment. However, reliable data about the IT infrastructure is needed to create a suitable cloud portfolio and to choose the right product variants, regions and parameters.

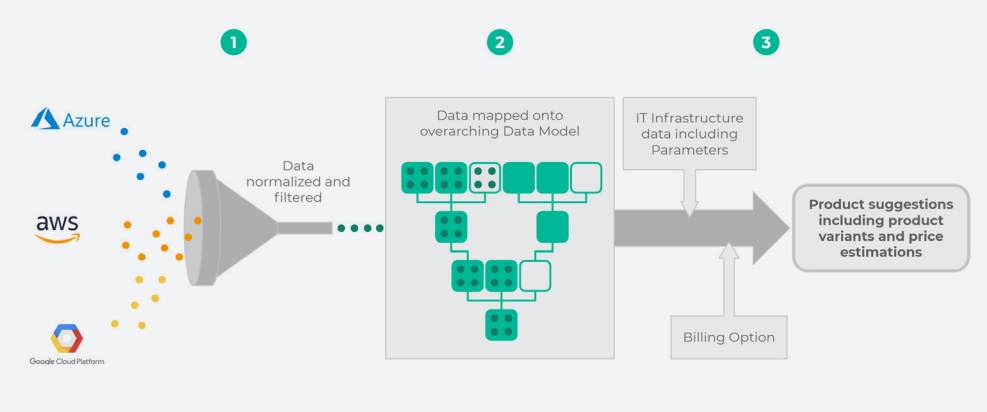

Figure 3: Txture solution visualized

At Txture we made it possible to automatically create cloud proposals and estimate prices based on IT infrastructure data: For each on-premises application and its underlying IT infrastructure cloud proposals are generated. Each proposal suggests a provider and the cloud products as well as the suitable product variants, suggested parameter settings like required storage or expected runtime per month and a location. In a next step, we can choose between different billing options. Information about the price options was previously collected together with the cloud prices via APIs. As a result, these proposals include a detailed overview of costs that we expect for the selected portfolio considering automatically selected parameters and the chosen billing option. Finally, parameters can be readjusted to refine the price estimates.

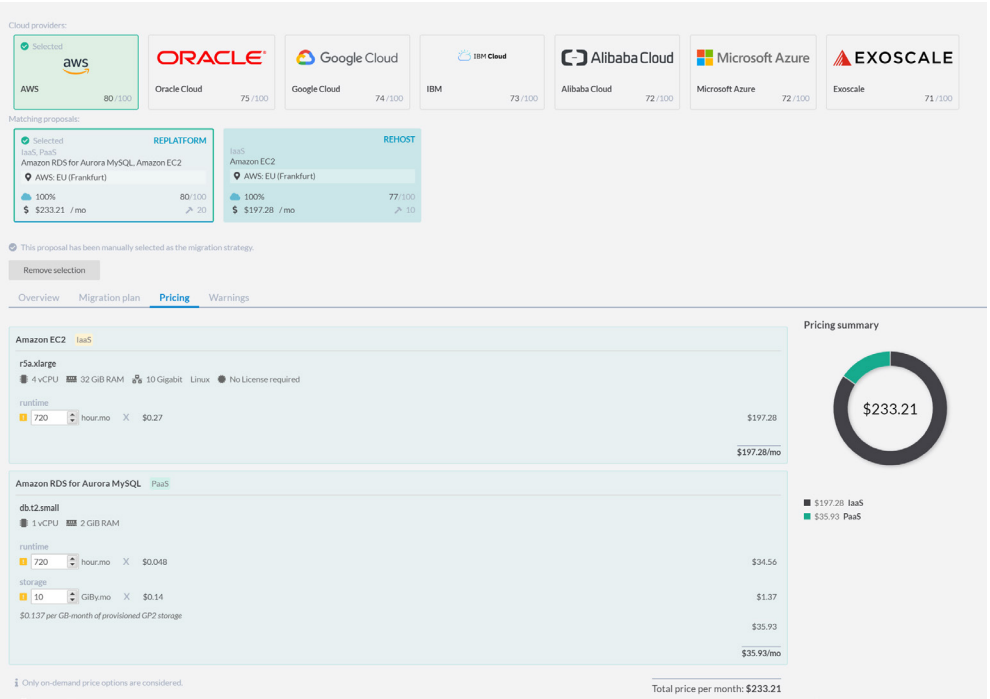

Figure 4: Example of a simple PaaS, IaaS deployment at AWS

It should be emphasized that this approach covers also difficult-to-compare PaaS products and comparisons between PaaS and IaaS options. Thus, for a single service, different service models can be compared. For example, if you choose to migrate your MySQL database to the cloud, you can compare the costs of an IaaS with an alternative PaaS solution.

Moreover, it is possible to compare different PaaS products like in the case of our MySQL database the two AWS products Amazon RDS (MySQL PaaS service) with Amazon Aurora (high-end MySQL PaaS service, including increased performance, backups, extended scalability, ...).

CONCLUSION

SURVIVING IN THE MESSY WORLD OF MULTI-CLOUD PRICING

1 Minutes

In the current situation on the cloud market makes it very difficult and time consuming to find a cloud portfolio that takes into account both technical aspects and cost-optimization. For large transformation projects, a manual or spreadsheet-based analysis of cloud costs is hardly feasible in the foreseeable future - especially when considering product portfolios of various cloud providers.

At Txture we have made it our mission to build an automated solution that helps companies with complex IT to make cloud prices tangible and predictable and to find the right cloud solution for their use cases. This can be an essential tool for the Cloud Center of Excellence.

For more detailed insights or discussions get in contact!

ABOUT THE AUTHOR
FLORIAN WIRTHENSOHN

Florian is an IT Transformation Consultant and Analyst at Txture. Besides his daily activities in pushing forward the customers' Cloud Transformation, Florian is passionate about researching promising topics around the future of IT. In addition to Cloud Transformations, Florian also specializes in Enterprise Architecture Management and IT Infrastructure Management.

florian.wirthensohn@txture.io

LinkedIn

+43 720 310 440 17

TXTURE

THE PLATFORM THAT GUIDES YOUR CCOE TO THE CLOUD

Txture is a software tool for members of the Cloud Center of Excellence (CCoE) of large organizations who want to speed up the cloud transformation and make the right migration decisions.

Our platform supports the end-to-end cloud transformation process from automated application landscape discovery, over cloud readiness assessment and roadmapping to hybrid cloud optimization. It also acts as a central cloud knowledge and decision hub.

Unlike other solutions, Txture takes the unique business, security, compliance and technical requirements of each business application into account, to maximize the business benefit of the migration, reduce transformation risk and save costs. All by making use of our Cloud Knowledge Engine that contains over 70.000 services of the hyperscalers such as Google, AWS, Azure or Alibaba.

With a constantly learning set of rules Txture analysis the cloud readiness of each application and proposes IaaS, CaaS, PaaS, and SaaS target architecture proposals based on the industry-specific transformation strategy of each organization. On this basis, large-scale cloud transformation is faster, safer and cheaper.

COMPANY

Txture GmbH, the company behind the innovative transformation platform, is a boutique software vendor with its headquarter in Innsbruck, Austria, that holds reselling and consulting partners across the globe.